

Towards an Information-Centric Delay-Tolerant Network

Gareth Tyson, John Bigham and Eliane Bodanese

School of Electronic Engineering and Computer Science, Queen Mary, University of London, UK

Abstract—Information-centric networks (ICNs) have emerged as a prominent future Internet architecture. They work on the principle that, in today’s society, a network should be optimised towards the delivery of information, rather than the transit of messages between pre-determined end hosts. At the same time, another prominent research direction has been the delay-tolerant networking initiative, which argues that networks must explicitly support the tolerance of high delays and disruptions in end-to-end paths. We believe that this observation is pervasive across a number of applications and network technologies and should therefore be strongly considered in any future ICN designs. This paper explores the potential of combining these concepts into an information-centric delay-tolerant network (ICDTN). Through a qualitative and quantitative investigation, we present arguments in favour of the design, as well as important research challenges.

I. INTRODUCTION

Information-centric networks (ICNs) [1][2][3] have blossomed into an exciting research topic, receiving attention from many researchers around the globe. Their key tenet is that our network infrastructure should centre around the effective delivery of content, rather than the delivery of packets between end hosts. Many compelling arguments have been made for its uptake, making it one of the most prominent future Internet proposals. However, it is evident that we are still in the early stages of its development. As such, it is important to expend significant effort to ensure that comprehensive network support is explicitly built into ICN designs from their inception.

Many important challenges have already been explored such as handling better security, scalability, network management and mobility. Despite this, there are clearly many other potential challenges that remain to be appropriately analysed and integrated into ICN architectures. One prominent example is the need for *delay/disruption tolerance* [4], a function that is increasingly important considering the emerging dominance of easily disrupted mobile communications [5]. Consequently, to ensure a fully integrated infrastructure, we posit that any future Internet designs should make delay tolerance a first-tier design consideration from the outset.

To this end, we explore the integration of information-centric and delay-tolerant principles to create an information-centric delay-tolerant network (ICDTN). Specifically, through analysis and simulations, we present arguments for an ICDTN, as well as identifying key opportunities and remaining challenges. As a use-case, we focus on the benefits of creating an ICDTN within a mobile human contact network [6]. However, we also extend our discussions into introducing delay-tolerant support into infrastructural (line-speed) ICNs.

The rest of the paper is structured as follows. After providing some background in Section II, we explore the key arguments for building an ICDTN in Section III, identifying the many similarities that exist between the two architectures. To help fuel discussion, we then formalise a straw man ICDTN model in Section IV before presenting our simulation study in Section V. Through these simulations, we explore the potential benefits (Section VI) that could be gained from combining delay-tolerant and information-centric principles, as well as discussing the key challenges that remain to be addressed (Section VII). Last, in Sections VIII and IX we present related work before concluding the paper.

II. BACKGROUND

There are two core aspects relating to this paper: information-centric networks (ICNs) and delay-tolerant networks (DTNs). In essence, an ICN is a network with the sole purpose of delivering content [1][2][3]. As such, an ICN exposes a publish/subscribe style abstraction unlike the existing Socket API [7]. This is because a host-centric network (HCN) is designed to scalably route packets from a source to a destination, whilst an ICN is designed to scalably deliver content from providers to consumers. As such, ICNs build routing [1] or resolution [2] mechanisms capable of passing requests to (optimal) sources of information.

In contrast, a DTN is a type of network that supports the existence of significant delays or disruptions between sending and receiving data [4]. Specifically, DTNs propose a store-and-forward architecture in which data units, termed *bundles*, can be temporarily stored at nodes (during network disruptions) until an appropriate next hop can be found [8]. These disruptions and delays could be caused by long distances (e.g. interstellar communications [9]) or, alternatively, by network partitions. The latter is of particular interest and can be caused by a number of events such as network failures, battery restrictions or wireless propagation limitations. One type of DTN that is particularly subject to the above types of events is the *pocket switched network* (PSN) [6]. PSNs are formed from opportunistic human contacts, typically by creating ad-hoc links between mobile phones [10]. These have found prominent applications such as message passing during cellular collapse in disaster situations [11], and constitute one of the key driving forces in the field of DTN. The rest of this paper now explores the potential of integrating all of the above principles into a single architecture that we term an information-centric delay-tolerant network (ICDTN).

III. ARGUMENTS FOR AN ICDTN

A number of benefits can be identified within both ICN and DTN architectures. An analysis of their characteristics is shown in Table I.¹ DTNs are an extremely flexible and resilient paradigm that offers huge benefits for unreliable or exotic networks that have previously relied on end-to-end connectivity. Their practicality in this regard has made them successful in many prototypes operating in various environments. However, most DTNs (i.e. RFC 5050 compliant [8]) rely on a host-centric routing mechanism, passing bundles between nodes regardless of the data being exchanged. In contrast, ICNs allow the network to gain a better understanding of the data itself, enabling it to be easily cached and reused. Such support offers a huge potential in disrupted environments that may rarely allow two hosts to reach other. Consequently, whilst delay-tolerant principles can offer *connectivity resilience* (e.g. through store-and-forward), information-centric principles can offer *information resilience* (e.g. through caching and replication). Interestingly, however, the key arguments for creating an ICDTN actually lie in their similarities rather than their differences. We identify four primary points of synergy between the two architectures:

- *Network Storage*: Both approaches rely on in-network storage.
- *Late Binding*: Both approaches espouse the late binding of names to locations.
- *Data Longevity*: Both approaches make the units of network interaction (ICN data objects, DTN bundles) long-term entities in comparison to traditional IP packets.
- *Flexible Routing*: Both approaches relax traditional constraints on routing and transport (loop free, uni-homed), making richer routing substrates possible.

These similarities seem to indicate that the integration of the two architectures is a logical progression, with the potential to combine the benefits of both. Importantly, we believe that the prominence of dis-connectivity in daily situations (e.g. underground transport, poorly provisioned areas, during expensive roaming) means that any future Internet architectures *must* place a high priority on supporting delay tolerance. Clearly, these examples motivate the need for mobile delay-tolerant support. However, we similarly argue that infrastructural ICNs should also include delay tolerance. On the one hand, this would allow better integration between infrastructural networks and more traditional mobile DTN environments. However, on the other hand, it would also open up a number of interesting possibilities. Specifically, the inclusion of delay-tolerant attributes in requests would allow superior resource allocation and scheduling within the network. This could, for example, allow delay-tolerant requests to be automatically performed during off-peak times to lower ISP economic costs [12]. Or, alternatively, it could allow new models of congestion control, by which networks could prioritise data that gets passed over particular links at given times (rather than having

intermediate routers dropping it). As such, we envisage that infrastructural ICNs could benefit from a top-down perspective (exploiting the delay tolerance of traffic), as well as from a bottom-up perspective (handling connectivity disruptions).

IV. AN ICDTN MODEL

Before continuing, it is important to formalise what we mean by an ICDTN. This section presents a model of the interactions that would occur in an ICDTN design. Our straw man ICDTN is constructed from a set of hosts that witness intermittent connectivity with each other. At any time, t , we assume a directed graph, $G(t)$, which consists of a set of nodes, $N = n_i$, and links, $L = l_{ij}$. Each node, n_i , possesses a set of Information Objects (IOs), $\mathcal{I}(n_i)$, as well as a set of Information Requests (IRs), $\mathcal{R}(n_i)$. An IO consists of an immutable piece of data, accompanied by a description (typically, this description would simply be a unique object identifier); whereas, an IR consists of criteria to match against an IO's description. Importantly, both IOs and IRs are persistent entities that can remain in the network via a store-and-forward mechanism until their TTL has been exceeded.

When a node, n_1 , generates an IR, a function, $l_select(ir)$, is used to select a set of active links, $\mathcal{T}(ir, n_1) \subseteq L(n_1)$, to forward the IR to, where $L(n_1)$ are the links of n_1 . Using flood-like routing, $|\mathcal{T}(ir, n_i)|$ would always equal $|L(n_i)|$, however, more intelligent strategies could also be considered in an attempt to minimise $\sum_{i=0}^{|N|} |\mathcal{T}(ir, n_i)|$. On receipt of the IR, a node, n_2 , computes the acceptability of the IOs in its local store, $\mathcal{I}(n_2)$, to create a new set, $\mathcal{K}(ir, n_2) \subseteq \mathcal{I}(n_2)$. Therefore, a function, $match(io, ir)$, exists that calculates the acceptability of an IO for a given IR using some metric, $m \in [0, 1]$. A threshold, τ , thus, defines the value required for an acceptable match; as previously mentioned, matching would often be based on unique identifiers resulting in $\tau = 1$. In such a situation, \mathcal{K} would be computed on n_i as follows:

$$\mathcal{K}(ir, n_i) = \{io | io \in \mathcal{I}(n_i) \wedge match(io, ir) = 1\}. \quad (1)$$

Once \mathcal{K} has been computed by n_2 , it is returned to n_1 . Following this, n_2 repeats the link selection process (l_select) to decide who the IR should be forwarded onto. Importantly, whenever a node, n_i , creates or receives an IR, it places it in its IR set, $\mathcal{R}(n_i)$, allowing it to be stored and passed to future nodes. As such, the link selection and forwarding process is repeated for each IR in $\mathcal{R}(n_i)$ whenever a state change occurs (i.e. a change to $L(n_i)$ or $\mathcal{I}(n_i)$). Due to this, responses may have to traverse multiple hosts; for example, if n_2 passes the IR to a node, n_3 , any responses to n_1 would have to return via n_2 . To exploit this property, transit IOs are also temporarily added to $\mathcal{I}(n_i)$ on any intermediate nodes (in this case n_2), allowing popular objects to be cached in the network.

V. EXPERIMENTAL METHODOLOGY

To study some of the quantitative benefits of integrating information-centric and delay-tolerant principles, we now focus on a single instantiation: an ICDTN operating in a human contact network (PSN). To this end, we have implemented

¹Due to space constraints, this table only presents a comparison between NDN [1] and the Bundle Protocol [8].

	ICN (NDN architecture)	DTN (Bundle Protocol)
Disruption Tolerance	Interest packets without FIB entries are dropped. Data packets without PIT entries are also dropped. Either case would occur during disruption	Bundles will be stored and forwarded during disruption, assuming sufficient storage capacity.
API	Synchronous get	Socket-like with chronological TTL setting
In-network storage	Ubiquitous router caching	Temporary store-and-forward buffer
Addressing/Naming	Hierarchical addresses for content	URI addresses for nodes
Security Model	Signing content by producers	Signing bundles by senders
Reliable Delivery	Receiver driven (re-request)	Bundle custody transfer and replication
Fragmentation	Managed by sources	Managed by network (and optionally source)

TABLE I
COMPARISON OF ICNS AND DTNS (USING NDN [1] AND BUNDLE PROTOCOL [8] ARCHITECTURES)

the above model of interaction into the ONE simulator [13]. Importantly, through this, we aim to identify key advantages and challenges within the field, rather than offering a comprehensive evaluation of an approach. Thus, a set of simulations have been performed using two modes: information-centric (ICN) and host-centric (HCN). Within the ICN mode, IOs are requested using their unique object identifiers as described in Section IV. To exploit this, each node also possesses an opportunistic cache that temporarily stores objects routed through it (default 20 objects). In contrast, within the HCN mode, requests are generated in an URL fashion based on predetermined locations, resulting in requests always being routed to the origin node that possesses the desired IO.

To gain meaningful results, we have taken a variety of mobility models to capture the links of nodes: (i) *INFOCOM'06*: contact traces taken from 98 nodes during the INFOCOM'06 conference (95 hours) [14]; (ii) *INFOCOM'05*: contact traces taken from 41 nodes during the INFOCOM'05 conference (76 hours) [15]; (iii) *NUS*: inferred contact patterns for 3000 students from the National University of Singapore (first 14 days) [16]; and (iv) *Helsinki*: 100 nodes performing random walks around ONE's Helsinki city map for 12 hours (speed uniformly between 30–60 metres per min and wait time 0–2 minutes). Using these nodes, we then generate IRs: periodically (every 1–5 minutes), a random node is selected to generate a request for an object. The IO to request is chosen from a 1k content set using a Zipf distribution. Interestingly, we observed little difference between different Zipf α parameters and, as such, set it to 0.8 throughout the experiments (in-line with previous observations [17]). Finally, IO sizes are uniformly distributed between 1KB and 4MB, and, at initiation, each IO resides on a single randomly selected host. To abstract away from any particular routing algorithm, we use epidemic routing in an attempt to understand the upper bounds of connectivity (epidemic has the highest delivery probability). This results in requests being flooded and then dropped once the request has been satisfied. This allows us to inspect whether a node can *theoretically* reach an IO without focussing on the underlying routing mechanism by which it is achieved.

VI. POTENTIAL BENEFITS

Through the above simulations, we now briefly explore some of the benefits and opportunities of integrating ICNs and DTNs (specifically PSNs due to the mobility models).

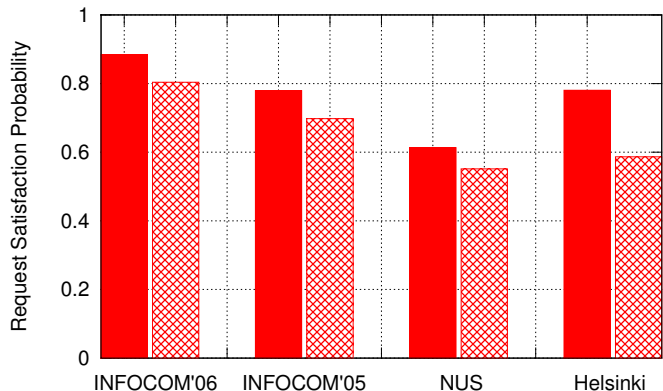


Fig. 1. Satisfaction Probability (solid is ICN, pattern is HCN)

A. Satisfaction Probability

We first inspect the probability by which IRs will be satisfied within their TTLs (each IR/IO has a TTL of 5 hours). Figure 1 presents the satisfaction probabilities (i.e. resilience). It can be seen that ICN outperforms HCN in all experiments, showing that increasing the number of endpoints capable of serving a node's needs (via caching and replication) does, indeed, result in a higher probability of requests being satisfied within a node's sphere of connectivity. Interestingly, we noticed little difference between smaller caches (20 objects) and larger caches (100 objects); for example, the INFOCOM'06 traces exhibited an improvement of less than 0.1% when increasing the cache size. Whilst this result is not surprising considering prior studies of ICN caching [18], it does highlight that significant resources are not necessary to achieve tangible benefits (an important factor on constrained devices).

A further observation is that each mobility model generates different results, despite the shared content request model. As expected, NUS performs worse because it has (temporally) sparser contact opportunities, whilst the INFOCOM traces perform the best due to their denser, smaller-scale nature. Evidently, there is no 'one-size-fits-all' optimal strategy; instead, some form of adaptive routing is likely to be needed.

B. Satisfaction Delay

Next, we look at the delay between nodes issuing requests and having them satisfied, presented in Figure 2 (this excludes

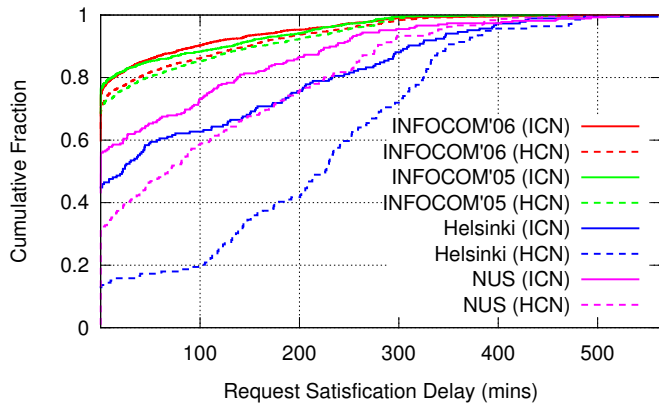


Fig. 2. Satisfaction Delay

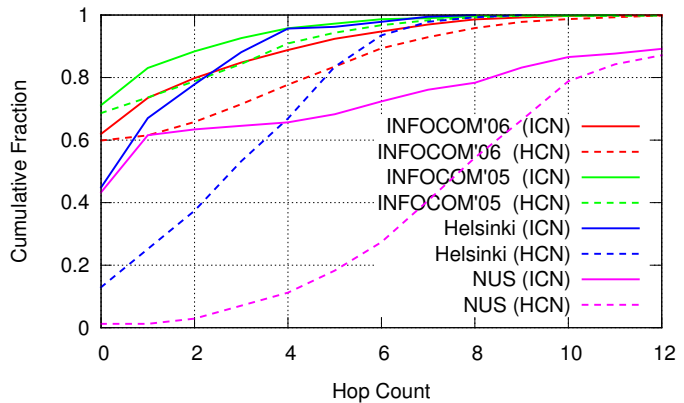


Fig. 3. Satisfaction Hop Count

requests that were not satisfied). It can be observed that the satisfaction delay of the ICN configuration is consistently lower than the HCN. Once again, this is due to the increase in the number of potential endpoints via caching, allowing nodes to retrieve content from nearby sources. This also has a positive impact on congestion by lowering the distance requests travel, which, in turn, further reduces queuing delays.

Interestingly, unlike the satisfaction probability, delay can also be improved noticeably by increasing cache sizes. Table II presents the delay when varying cache sizes; it can be seen that reductions of up to are $\approx 16\%$ can be achieved. As such, increasing cache sizes does not necessarily improve the ability to obtain an object but it does certainly improve delivery performance. It similarly creates potential for far more sophisticated routing algorithms, such as those that schedule IR propagation and IO storage based around the probability of IOs entering a neighbourhood in the (near) future. Clearly, this type of support is particularly important when considering flood-like routing, which has notoriously high overheads.

Size	INFOCOM'06		INFOCOM'05		NUS		Helsinki	
	Mins	Hops	Mins	Hops	Mins	Hops	Mins	Hops
20	25.4	1.2	27.3	0.7	67.9	6.6	101.9	1.3
100	21.2	1	25.1	0.3	64.2	4.8	88.7	1.2

TABLE II

AFFECT OF CACHE SIZE ON SATISFACTION DELAY (MINS) AND HOP COUNT

C. Hop Count

Finally, we look at the number of hops required for an IO to be returned to the requester, shown in Figure 3. This is related to the satisfaction delay, however, a key difference is that delay is heavily affected by inter-contact delays and queuing times, whereas hop count does not consider such underlying concerns. This can be best observed in the NUS trace, which shows quite different properties for delay and hop count measurements. The hop count therefore gives a more abstract metric; this is powerful because many things such as throughput and resilience can typically be modelled as monotonically decreasing functions of hop count. It can be seen that significant reductions can be achieved: in all

cases, the ICN mode results in fewer hops being required. Interestingly, a non-negligible number of requests can, in fact, be served from the local cache of the host. This is because popular objects rapidly become disseminated in the network (particularly in well connected networks, e.g. INFOCOM'05). This therefore could be better considered as a form of active replication, in which nodes receive objects before they have requested them. Clearly, this observation creates huge potential for more intelligent replacement algorithms that also take into account a node's personal interests when selecting objects to cache. Table II also highlights the ability to further reduce the hop counts via increased caching resources; for example, increasing cache sizes from 20 to 100 objects more than halves the hop count in the INFOCOM'05 traces due to the ease of IO dissemination in well connected networks. This can be compared against NUS, which serves very few local IOs because of the sparser connectivity, meaning IOs are just disseminated and cached within islands. Such observations obviously have potential to further feed into caching and routing decisions, e.g. adjusting local resource allocation based on the level of resilience required by the application.

VII. DISCUSSION AND REMAINING CHALLENGES

The above section has used some typical metrics to offer a preliminary analysis of benefits that could be achieved through an ICDTN compared against a DTN. However, we have also observed a number of interesting challenges.

A. Response Routing

Within traditional DTNs (and HCNs in general), the request and response routing process is symmetrical, i.e. they both route on the same algorithms and characteristics (host identifiers). This is because there is only a single network primitive: a bundle containing a source and destination. However, in an ICDTN, there are two primitives, requiring two routing mechanisms: (i) routing IRs to satisfactory IOs; and (ii) routing IOs back to requesting hosts. Clearly, the optimal processes by which these two tasks are achieved are different. To circumvent this, some approaches [1][19] rely on reverse path routing in which responses follow the reverse path of the

request. Unfortunately, however, this request-time binding of a response to a particular reverse path is infeasible in an ICDTN, which may suffer from frequent path changes. This means that it is necessary to maintain both information-centric and host-centric routing information, creating significantly greater load. A key challenge is therefore finding more robust ways to ensure that responses can maintain late-binding to paths, without requiring high levels of routing overhead.

B. Selective Cooperative Caching

Typically, on-path ICN caching is a simple process that exploits algorithms such as Least Recently Used. However, these principles work in infrastructural ICNs because packets are forwarded along deterministic paths. In contrast, most DTN routing protocols rely on stochastic *replication*, rather than deterministic *forwarding*. For example, epidemic routing involves the flooding of messages to all nodes, an approach that would theoretically homogenise every node's cache. As such, caching implicitly becomes a social process that must consider other nodes. It is therefore vital to re-think caching algorithms that apply to unicast environments. This is likely to involve greater intelligence, as well as cooperative algorithms.

C. Responder Selection

Within traditional DTNs (and HCNs in general), the selection of query responders is deterministic in that requests are addressed to pre-determined hosts. In infrastructural ICNs, the situation is similar as requests are typically unicast to specifically selected sources. However, within (opportunistic) ICDTNs, the selection of who should respond to a request is far more complicated because routing usually involves message replication, thereby allowing IRs to be received by many different sources. We noticed that this resulted in high levels of congestion due to IRs generating many responses. To address this within the simulations, we dropped requests once they had been satisfied; clearly, however, this would not be such a trivial task outside of a simulation environment.

Unfortunately, this is a notable challenge as most delay-tolerant routing protocols rely on some level of replicated multi-path forwarding. Further, the disconnected nature of the graph means that inter-node coordination is often impossible. As such, nodes must independently decide when to respond to requests. One possible approach to this is through passively observed popularity patterns (e.g. responding with an inverse probability to an IO's request probability). However, this is an area that still requires investigation.

D. Information Modelling

Currently, ICNs take a very simple approach to information modelling within the network. Each item of content receives a unique identifier, which can then be later used for retrieval. This is necessary for achieving line-speed switching, however, such approaches often rely on the ability to map human-understandable concepts onto the underlying identifier. Typically, this requires a resolution infrastructure; NetInf [3], for example, supports search functionality, allowing users to

map keyword queries to specific information identifiers. Such request/response functionality, however, is infeasible in a DTN environment in which connectivity is sparse. Consequently, it is important to embed as much information in the initial request as possible, mitigating the need for multiple round trips. A key challenge is therefore defining sufficiently expressive models to describe both information and queries. Importantly, such models must also assist the network in the previously discussed challenges (e.g. cooperative caching), as well as only requiring a low level of overhead.

E. Benefits for Infrastructural ICNs

The quantified benefits discussed in Section VI primarily relate to benefits that ICN principles can offer to mobile environments. ICN can be considered as a broad topic, however, currently, there are clear trends towards infrastructural line-speed deployments. As such, a further challenge is performing a quantitative analysis of the benefits of placing delay-tolerant support in such designs, e.g. adding chronological TTLs in NDN's Interest packets [1] to create a scheduling-like service within the network. Some potential applications of this have been discussed in Section III, however, these must be appropriately evaluated. Alongside this, it is similarly vital to investigate the interconnect between any mobile ICDTNs and infrastructural ICDTNs. Generally, DTNs connect to infrastructural (IP) networks via a gateway, however, an important element of any future Internet architecture is the ability to unify these different network components without the need for complicated translations. As such, a further key research challenge is to design technologies and protocols that can offer 'plug-and-play' behaviour between the different incarnations of ICNs (this similarly applies to different ICN protocols).

VIII. RELATED WORK

Within the paper, we have attempted to avoid the specific details of any particular system, instead, performing a more generalised analysis. However, there are a number of related ICNs and DTNs. Currently, most ICNs focus on infrastructural Internet-scale delivery. As such, their routing and addressing targets scalability, typically using simple names. Prominent examples include NDN [1], NetInf [3], PURSUIT [2] and Juno [20]. These all uniquely identify information objects rather than hosts, offering the ability to discover and retrieve them when desired. Although, there is certainly potential to extend these into the mobile and delay-tolerant domain, as of yet, less focus has been given to this (discussion of this can be found in [21]). Having said that, some independent information-centric routing solutions have also been proposed (e.g. LFBL [22], Slinky [23]), supporting ad hoc communications.

Alongside the above ICN-themed work, there has also been a range of research into DTNs. This has been led by efforts that have resulted in standards such as RFC 5050 [8], based around the architecture originally proposed by Fall [4]. A key driving scenario behind this work has been the Interplanetary Internet [9], in which distances create significant communication delays. Alternatively, in terms of the discussed

usage domain, pocket switched networks [6] are most related; specifically, those which perform opportunistic content sharing. These include architectural designs such as Hagggle [10], as well as a range of delay-tolerant routing proposals, e.g. B-SUB [24] and SocialCast [25]. Systems such as Hagggle, SPINDLE [26] and ContentPlace [27] are based around the concept of opportunistic content exchanges. In Hagggle, for instance, when two nodes meet, they exchange queries (called node descriptions) and then exchange any data objects that match the queries. Importantly, in some of these systems, their content-oriented nature makes them implicitly delay-tolerant, allowing content routing to be performed over time as well as space. Alongside this, efforts have also been made to integrate information-centricity into the existing DTN standards. The Bundle Protocol Query (BPQ) [28], for example, allows bundles to be augmented with query information that stipulates bundles of interest, allowing on-path replicas to be returned.

Finally, as well as mobile and interplanetary DTNs, there is also work looking at delay tolerance in traditional fixed infrastructure. Laoutaris *et al.* [12] analysed the potential of time-shifting delay-tolerant traffic to avoid peak times within the Internet. Similarly, Chhabra *et al.* [29] looked at offering incentives for home users to do this. We believe these papers offer a further powerful use-case for ICDTNs, which would allow such mechanisms to be realised as native network operations. This would similarly make a powerful business case for ISP uptake, considering the economic benefits possible [12].

IX. CONCLUSION AND RESEARCH AGENDA

This paper has explored the potential of integrating information-centric and delay-tolerant principles into a shared ICDTN architecture. Our preliminary experiments show that there are notable benefits over traditional PSNs/DTNs, namely in terms of request satisfaction probability, satisfaction delay and hop count. When placed in the context of our increasingly mobile world [5], it is therefore undeniable that there is significant potential in the area. However, similarly, there remain a number of challenges (which also constitute our future work). Specifically, we believe efforts should be expended on (i) *decentralised matching algorithms*: strategies that allow nodes to generate sophisticated, rich queries, rather than relying on ICN's current more simplistic name-based querying; (ii) *response management*: cooperative algorithms that allow replicas and sources to coordinate responses to ensure that the minimal number of optimal source(s) respond to avoid saturating the limited network resources; and (iii) *integration*: internetworking technologies that allow the seamless integration of mobile ICDTNs with (multiple) wider infrastructural ICNs without the need for complex bridging and translation. This last point, similarly, relates to the need to appropriately analyse the feasibility and benefits of deploying delay-tolerant functionality into these infrastructural ICNs.

REFERENCES

[1] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," in *In Proc. 5th ACM CoNEXT*, 2009.

[2] D. Trossen and G. Parisi, "Designing and realizing an information-centric internet," *IEEE Communications Magazine*, vol. 50, july 2012.

[3] D. Kutscher, S. Farrell, and E. Davies, "The NetInf Protocol." Internet Draft, 2012.

[4] K. Fall, "A delay-tolerant network architecture for challenged internets," in *Proc. SIGCOMM*, 2003.

[5] "Cisco visual networking index: Forecast and methodology, 2011-2016."

[6] N. Sastry, D. Manjunath, K. R. Sollins, and J. Crowcroft, "Data delivery properties of human contact networks," *IEEE Trans. Mob. Comput.*, pp. 868–880, 2011.

[7] G. Tyson, A. Mauthe, S. Kaune, P. Grace, and T. Plagemann, "Juno: An adptive delivery-centric middleware," in *Proc. 4th Intl. Workshop on Future Media Networking (FMN)*, 2012.

[8] K. Fall and S. Burleigh, "Bundle Protocol Specification." RFC 5206 (Experimental), Nov. 2007.

[9] S. Burleigh, A. Hooke, L. Torgerson, K. Fall, V. Cerf, B. Durst, K. Scott, and H. Weiss, "Delay-tolerant networking: an approach to interplanetary internet," *Communications Magazine, IEEE*, vol. 41, pp. 128 – 136, june 2003.

[10] J. Scott, P. Hui, J. Crowcroft, and C. Diot, "Hagggle: a Networking Architecture Designed Around Mobile Users," in *IFIP WONS*, 2006.

[11] P. Jiang, J. Bigham, E. Bodanese, and E. Claudel, "Publish/subscribe delay-tolerant message-oriented middleware for resilient communication," *IEEE Communications Magazine*, vol. 49, september 2011.

[12] N. Laoutaris, G. Smaragdakis, P. Rodriguez, and R. Sundaram, "Delay tolerant bulk data transfers on the internet," in *Proc 11th SIGMETRICS*, 2009.

[13] A. Keränen, J. Ott, and T. Kärkkäinen, "The one simulator for dtn protocol evaluation," in *Proceedings of the 2nd International Conference on Simulation Tools and Techniques*, 2009.

[14] J. Scott, R. Gass, J. Crowcroft, P. Hui, C. Diot, and A. Chaintreau, "CRAWDAD trace cambridge/hagggle/imote/infocom (v. 06-01-31)."

[15] J. Scott, R. Gass, J. Crowcroft, P. Hui, C. Diot, and A. Chaintreau, "CRAWDAD trace cambridge/hagggle/imote/infocom2006 (v. 09-05-29)."

[16] V. Srinivasan, M. Motani, and W. T. Ooi, "CRAWDAD data set nus/contact (v. 06-08-01)."

[17] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and zipf-like distributions: Evidence and implications," in *Proc. 18th IEEE INFOCOM*, 1999.

[18] G. Tyson, S. Kaune, S. Miles, Y. El-Khatib, A. Mauthe, and A. Taweel, "A trace-driven analysis of caching in content-centric networks," in *Proc. Intl. Conference on Computer Communication Networks (ICCCN)*, 2012.

[19] M. G. Soon-Young Oh, Davide Lau, "Content centric networking in tactical and emergency MANETs," in *Proc. IFIP Wireless Days Conference*, 2010.

[20] G. Tyson, A. Mauthe, S. Kaune, P. Grace, A. Taweel, and T. Plagemann, "Juno: A middleware platform for supporting delivery-centric applications," *ACM Transactions on Internet Technology*, 2012.

[21] G. Tyson, N. Sastry, I. Rimal, R. Cuevas, and A. Mauthe, "A survey of mobility in information-centric networks: challenges and research directions," in *Proc. 1st ACM MobiHoc NoM Workshop*, 2012.

[22] M. Meisel, V. Pappas, and L. Zhang, "Ad hoc networking via named data," in *Proc. MobiArch*, 2010.

[23] V. Kawadia, N. Riga, J. Oppen, and D. Sampath, "Slinky: An adaptive protocol for content access in disruption-tolerant ad hoc networks," in *Proc. MobiHoc Workshop on Tactical Mobile Ad Hoc Networking*, 2011.

[24] Y. Zhao and J. Wu, "B-sub: A practical bloom-filter-based publish-subscribe system for human networks," in *Proc 30th ICDCS*, 2010.

[25] P. Costa, C. Mascolo, M. Musolesi, and G. Picco, "Socially-aware routing for publish-subscribe in delay-tolerant mobile ad hoc networks," *IEEE Journal on Selected Areas in Communications*, vol. 26, june 2008.

[26] R. Krishnan *et al.*, "The spindle disruption-tolerant networking system," in *Proc. IEEE MILCOM*, 2007.

[27] C. Boldrini, M. Conti, and A. Passarella, "Contentplace: social-aware data dissemination in opportunistic networks," in *Proceedings of the 11th international symposium on Modeling, analysis and simulation of wireless and mobile systems*, 2008.

[28] S. Farrell, A. Lynch, D. Kutscher, and A. Lindgren, "Bundle protocol query extension block." draft-farrell-dtnrg-bpq-01, Mar. 2012.

[29] P. Chhabra *et al.*, "Home is where the (fast) Internet is: flat-rate compatible incentives for reducing peak load," in *Proceedings of ACM SIGCOMM HomeNet workshop*, 2010.